Monitoring high-speed telecom networks

# 10Gbps and beyond

Polystar

## Polystar OSIX AB

www.polystar.com/OSIX

+46 8 50 600 600

Polystar OSIX offers the telecommunications market solutions for maximising network performance, as well as tools helping operators to provide the best possible customer experience.

With the Polystar portfolio of monitoring and analysis tools, an operator can meet market change and focus sharply on service and customer experience. Today, many multinational telecom operators use our Customer Assurance and Service Assurance solutions to get real-time visibility of network quality and customer experience—with the ultimate aim of maximising long-term customer loyalty.

Founded in Stockholm in 1983, Polystar has grown to a global presence with over 100 customers in more than 50 countries. The company has always been profitable, and last year showing a remarkable 40% year over year growth.

# 10Gbps and beyond

Jonas Lesser, Polystar OSIX AB, July 2010

This article describes the problems of monitoring modern 10Gbps telecom networks and some of the solutions. We will also briefly describe the custom-built system Polystar OSIX is using to tackle these issues.

The Polystar OSIX 10GbE MediaProbe uses an FPGA-based MicroTCA system as the interface to the monitored network. Tailor-made algorithms in the FPGA have been developed based on studies of real-world scenarios.

Mobile phone networks today experience a tremendous increase in volume, not only due to an increased number of customers but also owing to the new ways these customers use the mobile connections. The growth in traffic, driven by smartphones and IP-based services like video chats, YouTube, movie streaming, and IPTV, is putting extreme pressure on the communications infrastructure to deliver more bandwidth. Network operators are starting to deploy 10Gbps network infrastructure to handle the increased bandwidth requirements.

Mobile phone network operators are now also internet service providers, and consequently need new tools to inspect and monitor their networks and packet-based services.

Customers' demand for increased bandwidth and high connection success implies that operators need to analyse a vast amount of data in order to have sufficient information should quality or support issues arise. These customer demands and support issues drive the fact that the new tools must be able also to handle the user data, not only the control signalling.

Regardless of architectural approach to the monitoring infrastructure, each processor node in the monitoring infrastructure can only handle a limited amount of data, depending on the amount of computing power available to a specific application. When traffic increases over this level, all data cannot be processed by one monitoring processing node, so the payload must be processed on several processors in the monitoring system. Messages belonging to the same user session must, however, be processed by the same unit to be able to extract reliable information for the session.

Minimising data by filtering out control signalling data and discarding user data is no longer an option, since the actual user data now is as important as the control signalling to assess the session quality. In addition, some protocols use in-band signalling, mixing control messages with the user data stream. This means that all data must be inspected to find the control signals.

Intelligent load sharing based on information from deep packet inspection is necessary to enable the nodes to use the algorithms suited for each problem.

In some cases, it is suitable to use the source-destination IP address pair to split the data over several monitoring processing nodes. Sometimes, such an approach is not sufficient, for example when a network element handles more traffic than a single processing node can handle. When this is the case, other header fields can be used. If the monitored data is tunnelled, as for example on Gn, the tunnelling identifier can be used. This will still keep the user sessions at the same processor node. Packets that are control data or charging data can be selected using header inspection and handled separately to enable special user features, such as real-time session control analysis.

Using intelligent load-sharing algorithms also means that the management can be kept to a minimum. Instead of using manually inserted load-sharing rules each time the monitored network changes, the algorithms are kept and remain stable even after reconfiguration.

Time stamping is important to be able to decide the order in which probed packets are recorded by the probe-system. A correct and sufficiently accurate time stamping can align data arriving from multiple probes or different regions of the network. Time stamping can also be important if the quality of the data streams has to be measured, specifically if jitter measurements are needed.

As the bandwidth increases, the precision of the time stamping also needs to increase. At 10Gbps, time stamping in the CPU of a server is not good enough. Incoming packets must be time stamped before they reach a buffer to ensure a deterministic behaviour.

In the 10GbE MediaProbe, Polystar OSIX uses an FPGA to time-stamp the packets with nanosecond resolution.

In some instances, the traffic at a given point in the network can be split over several interfaces. The split can be controlled by a range of parameters or simply use a spill-over behaviour. Often, the monitoring equipment needs to process the data using a different division than the split used by the network. If the first unit of the monitoring equipment has the ability to forward the data in a specific and predictable way, all monitored points can be used as if the data was collected from just one access point. In addition, the accurate time stamp can be kept across the multiple interfaces. The Polystar OSIX 10GbE monitoring system can be used in such a way.

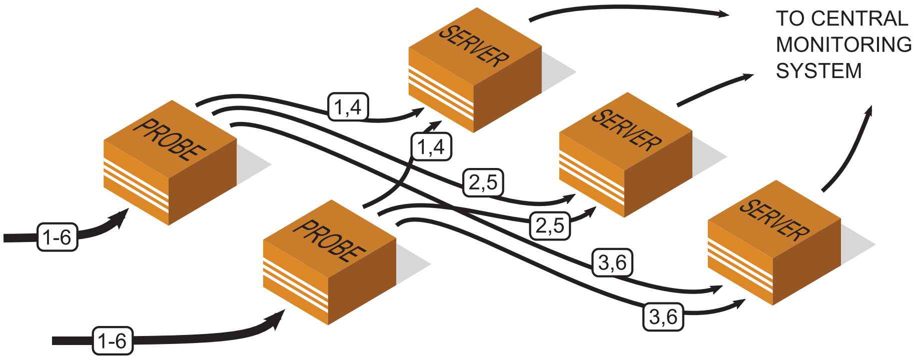


Figure 1. Traffic split over two incoming lines are load-shared to several probe servers. The same load-share algorithm is used on both probes, keeping each session confined to one server.

When investigating traffic on the backbone of several operators, we found that the packets were often fragmented. In some cases, the fragmentation was such that it hindered the monitoring equipment from correctly gathering information. The fragmentation is caused by certain configurations of the core equipment of the operator. To fully support all scenarios, Polystar OSIX's monitoring system uses a custom reassembly algorithm to take care of such cases. The reassembled packets can then be load-shared to the correct destination.

When a specific type of traffic needs to be monitored more carefully, pattern-matching algorithms can also be inserted into the FPGA of the monitoring system. These can be used to tag and direct certain packets in a fully saturated 10Gbps line, and help the subsequent monitoring units to collect statistical data. This can increase the efficiency of the monitoring system and minimise the amount of hardware required for the task.

Although it might in some cases be possible to accommodate special cases of 10GbE traffic using standard servers with offloading NICs and optimised code, the introduction of a scalable and high performance topology based on special hardware facilitates later system migration, once 40GbE and 100GbE become a reality. In addition, this solution will be much more flexible and easy to manage when the monitored network evolves.

The Polystar 10GbE monitoring system is built using components that have well-defined roadmaps for 40Gbps and 100Gbps throughput. When there is a need for 40Gbps or 100Gbps monitoring, the individual units can be upgraded without the need to modify the overall system structure.

For high loads on multiple 10GbE interfaces, a hardware-based tailor-made solution is the only option.

## About the author

During his PhD studies at Stockholm University, Jonas Lesser worked on a data digitalisation and readout system for the ATLAS detector of the LHC particle accelerator at CERN.

Since 2005, he has been System Manager for Polystar's hardware development, where his work included the Gemini system for STM-1 links. He has also designed the new 10Gbps monitoring system for the MediaProbe.

To his Polystar colleagues, he is also known as the bass player in the company rock band Polystarz.